

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

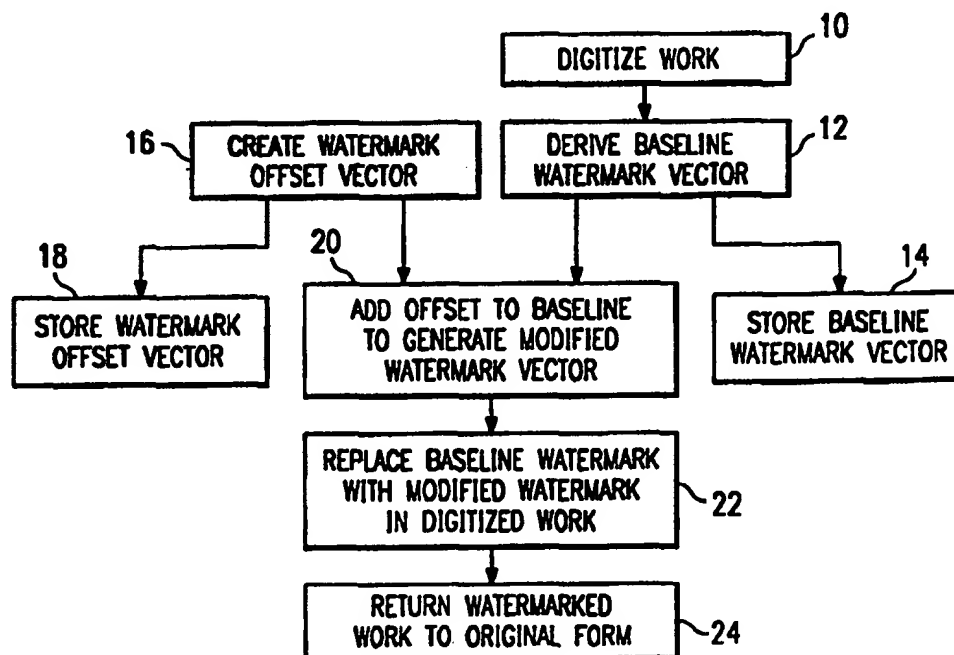
**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G09C 3/06	A2	(11) International Publication Number: WO 99/10858 (43) International Publication Date: 4 March 1999 (04.03.99)
(21) International Application Number: PCT/US98/17833 (22) International Filing Date: 28 August 1998 (28.08.98) (30) Priority Data: 08/920,365 29 August 1997 (29.08.97) US (71)(72) Applicant and Inventor: LEIGHTON, F., Thomson [US/US]; 15 Charlesden Park, Newtonville, MA 02160 (US). (74) Agent: JUDSON, David, H.; Hughes & Luce, L.L.P., Suite 2800, 1717 Main Street, Dallas, TX 75201 (US).		(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>Without international search report and to be republished upon receipt of that report.</i>

(54) Title: METHOD FOR PROTECTING CONTENT USING WATERMARKING



(57) Abstract

A watermarking procedure wherein each of a set of copies of the work has a slightly-modified form of a "baseline" watermark that is placed within a critical region of the data. The slight variations in the watermarks, however, are not perceptually visible and do not interfere with the work. If multiple persons collude to attempt to create an "illicit" copy of the work (i.e., a copy without a watermark), however, at least one of the modified watermarks is present in the copy, thereby identifying both the illicit copy and the copier.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

Method For Protecting Content Using Watermarking

This application is a continuation-in-part of prior co-pending application U.S. Serial No. 08/615,534, filed March 12, 1996, now U.S. Patent
5 No. 5,664,018.

TECHNICAL FIELD

The present invention relates generally to preventing unlawful copying of audio, video and other media that can be digitized and, more particularly, to
10 improved watermarking techniques that are robust even against multiple individuals who conspire together with independent copies.

BACKGROUND OF THE INVENTION

The proliferation of digitized media (audio, image and video) and the ease with which digital files can be copied has created a need for copyright
15 enforcement schemes. Conventional cryptographic systems permit only valid keyholders access to encrypted data, but once such data is decrypted there is no way to track its reproduction or retransmission. Such schemes thus provide insufficient protection against unauthorized reproduction of information. It is known in the prior art to provide a so-called digital "watermark" on a document
20 to address this problem. A "watermark" is a visible or preferably invisible identification code that is permanently embedded in the data and thus remains present within the data after any decryption process. One example of a digital watermark would be a visible "seal" placed over an image to identify the copyright owner. However, the watermark might also contain additional
25 information, including the identity of the purchaser of a particular copy of the material.

Many schemes have been proposed for watermarking digital data. In a known watermarking procedure, each copy of a document D is varied slightly so as to look the same to the user but also so as to include the identity of the
30 purchaser. The watermark consists of the variations that are unique to each copy. The idea behind such schemes is that the watermark should be hard to

remove without destroying the document. Thus, a copy of a watermarked document should be traceable back to the specific version of the original from which it was created.

Although many prior art schemes claim to possess the "unremovable" property, all existing schemes are easily defeated by the following type of attack. Assume the attacker obtains two copies of the document that is being protected by the watermarking scheme. Each copy may have a different watermark, neither of which is supposed to be removable. The attacker now makes a third version of the document (which he hopes will not have a traceable watermark) by averaging his two copies. For a pictorial document, for example, each pixel of the third version would be the average of the corresponding pixels in the watermarked copies.

Using existing approaches to watermarking, the third copy of the document produced by the attacker will look like the original versions but the watermark will be destroyed. This is because the "average" of two watermarks does not carry sufficient information to be tied to either of the watermarks individually. Thus, the watermarking scheme can be rendered ineffective by simply averaging two copies of the document.

There is thus a need to devise a watermarking scheme that is immune to these and other such attacks, especially those in which the adversary obtains multiple copies of the original document.

BRIEF SUMMARY OF THE INVENTION

It is the principal object of the invention to describe a digital watermarking scheme wherein the watermark is robust against collusion by multiple individuals who each possess a watermarked copy of the data.

It is another object to describe such a scheme wherein the watermark cannot be removed by an adversary who obtains multiple copies of the original work.

It is a more general object of the invention to describe a watermarking method that is secure against any form of attack including, without limitation, averaging attacks.

It is still a further object of the invention to describe a watermarking
5 procedure wherein each of a set of copies of the work has a slightly-modified form of a "baseline" watermark that is placed within a critical region of the data. The slight variations in the watermarks, however, are not perceptually visible and do not interfere with the works. If multiple persons collude to attempt to create an "illicit" copy of the work (i.e., a copy without a watermark), however,
10 at least one of the modified watermarks is present in the copy, thereby identifying both the illicit copy and the copier.

It is still thus another object to describe a watermarking scheme of the type recited above wherein combining copies of the same data set does not destroy the watermark.

15 It is a further object of the invention to describe such a watermarking scheme that may be used to identify one or more of the parties who are colluding to destroy the watermark.

It is another more general object of the invention to describe a digital watermarking process that may be used as evidence in a Court because it is
20 robust against collusion.

According to the preferred embodiment of the invention, the work to be protected is digitized into a data file or string of data. A first digital watermark is then inserted in a first copy of the data file, preferably in a critical region of the data. A "critical" region may consist of the entire document or alternatively
25 will be some valuable portion of the work that will end up being significantly corrupted if the watermark is corrupted. A second digital watermark is then inserted in a second copy of the data file in a similar manner, and the process is repeated for additional copies. According to the invention, the first and second digital watermarks are slight variations of a "baseline" watermark,
30 which is kept secret, and one cannot perceive any differences between the first

and second copies due to these variations. The baseline watermark may be a digital string that is part of the original data being protected. Preferably, the variations are "randomized" in such a manner that if two persons were to collude to attempt to create an "illicit" copy of the work (i.e., a copy without a watermark), at least one of the first or second watermarks would still be present in the copy. After the watermark is inserted into the work, the work can be converted back to its original form.

Thus, the scheme ensures that different possessors of watermarked copies of a work cannot create a "clean" copy that does not include at least one of the slightly-modified watermarks. Indeed, by comparing the watermark of the illicit copy with the baseline watermark, one can determine the identity of the forger.

Although not meant to be limiting, preferably the "variations" are generated using a "random" offset, and in particular a "normal distribution."

15 BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and the advantages thereof, reference should be made to the following Detailed Description taken in connection with the accompanying drawings in which:

FIGURE 1 is a block diagram illustrating the method of inserting a digital watermark into a copy; and

FIGURE 2 is a block diagram illustrating the method for retrieving a digital watermark from a copy and correlating the retrieved watermark with a stored watermark.

DETAILED DESCRIPTION

25 According to the invention, the work to be protected may be an image (photographs and graphics), video and/or audio (speech and music). The particular type of work is not relevant to the invention. Referring now to FIGURE 1, the work, in whatever form, is digitized at step 10 into a data file or string of data either as part of the inventive technique or through some known A/D preprocessing. In the invention, there is a "baseline" watermark that is

preferably stored and not used in making a particular copy of the work (although this step is not necessarily required). This baseline watermark is then processed to create a set of one or more "modified" watermarks, each of which is related to the baseline watermark in a predetermined manner.

5 Preferably, the "offsets" needed to create the modified watermarks are not fixed, however, but are "randomized." In this way, a very small amount of "noise" is added to the offsets that does not alter the perception of the watermarked copies but still ensures that possessors of such copies cannot collude to remove all existence of the watermark in at least one illicit copy.

10 In general, collusion-type attacks are prevented according to the invention by constructing a watermark using randomness in a specific way. Preferably, an n -length digital string: x_1, x_2, \dots, x_n is derived at step 12 from the data to be watermarked and stored at step 14 for future reference. This may be referred to as the "baseline" watermark. The string is preferably "critical" to
15 the data in that corruption of the string will corrupt the data in a way that can be perceived and which will diminish the value of the corrupted document. Generation of the baseline watermark can be achieved in many ways, e.g., by digitizing some portion of the document and using the resulting data or some subset thereof. (Whatever method is used is also used in the verification
20 process, as discussed below). An n -length watermark vector w_1, w_2, \dots, w_n is then created at step 16 and stored at step 18 for future reference. The vector is preferably created by choosing each w_i from a specified random distribution (preferably the normal distribution). The random distribution used for each w_i may or may not be the same (e.g., depending on whether it is desired to embed
25 some specific serial number data in the watermark). The watermark vector is then added at step 20 to the string x_1, x_2, \dots, x_n , and the result reinserted at step 22 into the original data to be protected. The work may then be converted back to its original form (image, video, audio, etc.) at step 24.

Assume it is now desired to retrieve the watermark from a copy D' . This
30 can be accomplished, as shown generally in FIGURE 2, by digitizing the copy

D' at step 30 and then computing at step 32 the derived values x_1', x_2', \dots, x_n' , using the same algorithm used to compute the baseline watermark. Then, the method proceeds at step 34 by retrieving the original base line watermark, x_1, x_2, \dots, x_n , from memory and subtracting out x_1, x_2, \dots, x_n from x_1', x_2', \dots, x_n' to

5 compute a derived watermark w_1', w_2', \dots, w_n' at step 36. A correlation value (preferably an inner product) is then calculated between the derived watermark and w_1, w_2, \dots, w_n , retrieved at step 38, to produce a correlation value at step 40. The correlation value is compared at step 42 to threshold levels, and if the correlation is high (step 44), then there is a match and a watermark is present.

10 If the correlation is low (step 46), the watermark is not present. (The inner product scheme works by computing the absolute value of the sum $w_1 w_1' + \dots + w_n w_n'$).

This scheme is immune to collusion because the watermark is random and because different watermarks are completely uncorrelated. In existing

15 schemes, different watermarks are highly correlated and so it is easy for an attacker to exploit the correlation to destroy the watermark (e.g., by an averaging attack). In the invention method, there is simply not enough information contained in "t" different watermarked copies of the data in order for the adversary to remove the watermark. More specifically, if the attacker

20 obtains "t" copies of watermarked data using the normal distribution to construct the watermarks (with watermarks w_{11}, \dots, w_{1n} , through w_{t1}, \dots, w_{tn}), it will appear to the attacker as if the original baseline watermark is $x_1 + (w_{11} + \dots + w_{t1})/t, \dots, x_n + (w_{1n} + \dots + w_{tn})/t$, which is not the true baseline watermark x_1, \dots, x_n . The distinction is important since the former string is correlated with each

25 of the watermarks $w_{11} \dots w_{1n}$ through $w_{t1} \dots w_{tn}$. In other words, the attacker simply does not have enough information in order to evade the watermark, no matter what sort of attack is used. Hence, one can prove that either the attacker must destroy the data or he must leave a trace of at least one of the component watermarks which will be revealed when the correlation test is run.

Only someone with knowledge of the original baseline watermark could remove the watermark without detection.

Therefore, "m" copies of the work include variations of a baseline watermark such that up to "t" persons who possess those copies cannot
5 collude to create a "clean" copy (i.e., one without any watermark whatsoever). Stated another way, any "t" persons who collude in such a manner will always create an illicit copy that includes one of the modified watermarks. Comparison of the watermark of the illicit copy with the baseline watermark then identifies what party made the copy (assuming there is a record of which party originally
10 got which "version").

According to a preferred method, a first digital watermark is inserted in a first copy of a data file, preferably in a critical region of the data. A second digital watermark is then inserted in a second copy of the data file in a similar manner, and the process is repeated for additional copies. As discussed
15 above, the first and second digital watermarks are slight variations of a "baseline" watermark, which is kept secret, and one cannot perceive any differences between the first and second copies due to these variations. Preferably, the variations are "randomized" in such a manner that if two persons were to collude to attempt to create an "illicit" copy of the work (i.e., a
20 copy without a watermark), at least one of the first or second watermarks would still be present in the copy. In the preferred embodiment, a watermark consists of a sequence of numbers $W = w_1, \dots, w_n$, where each value w_i is chosen independently and approximately according to $N(0,1)$ (where $N(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2). The watermark may
25 consist of a number (e.g., 1000) of randomly generated numbers with a normal distribution having zero mean and unity variance. Alternatively, w_i could be selected according to $N(\mu_i, \sigma_i)$ where μ_1, \dots, μ_n can be a serial number corresponding to the copy being watermarked (or other information that may be embedded).

In order to detect the presence of a watermark W in a derived watermark signal W' , we preferably use a correlation function $\text{cor}(W, W') = |W \bullet W'|$, which is the inner product of two vectors. If W were selected according to the normal $N(0,1)$ distribution and W' is uncorrelated to W (but of the same order), then
5 the correlation will be small (about \sqrt{n}). If W' is closely correlated to W , then the correlation will be large (about n). If W' is uncorrelated to W but is of a larger order (e.g., due to intentional or unintentional noise or attempts to hide the watermark), then the correlation might also be large. (Specifically, if W' is uncorrelated to W but has B times the magnitude, then the correlation is about
10 $B \sqrt{n}$. If B is large, then the data D' will not resemble D . (The notion of large in this context depends on the application and the level of security/clarity desired). In any event, the watermark is said to be present if $\text{cor}(W, W') > c \sqrt{n}$, where c is a predetermined constant that depends on the application and level of security desired (e.g., $c=4$).

15 The correlation will be low if the watermark is not present and the work is not destroyed. The correlation will be high if D' was derived from the watermarked document or if the data has been corrupted beyond recognition (the latter condition being determined by inspection).

As noted above, it is preferable that each of the "modified" watermarks
20 be placed in a critical region of the data. Of course, the exact location will depend on the nature of the work being protected. It is also helpful if every entry in this region of data is largely uncorrelated with the other data. It has been suggested (by Cox et al) that this can be accomplished by embedding a watermark in the spectrum of an image, the temporal frequency domain of an
25 audio signal, or the spatio-temporal frequency domain of a video sequence. Although the above techniques are preferred, one may even encode the watermark in other less, desirable places (such as in the low order or least significant bits) of the data and still obtain the advantages of the collusion-resistant feature of the invention where multiple parties may collude to remove
30 the watermark.

Variations

In the embodiment discussed above, the original document (or an original baseline watermark vector) is stored in order to determine whether the watermark is present in a copy of the document. In the embodiment previously
5 described, the original baseline watermark vector is retrieved at step 34 and subtracted from the derived baseline watermark vector to produce the derived offset watermark vector. This step can be omitted without changing the detection protocol or its results. In particular, the derived offset watermark vector may be set equal to the derived baseline watermark vector. This
10 change increases the noise level in the correlation test, but not beyond tolerable levels. Further, the noise levels can be reduced by specially selecting the original offset watermark vectors to have low noise (e.g., by selecting them to be orthogonal to the original baseline watermark vector to which they are being applied) or by running the correlation test on only specific
15 components of the vectors.

Another improvement would be to remove the need to store the original offset watermark vector. As discussed above, in one embodiment of the invention it is necessary to store a copy of the original offset watermark vectors (see, e.g., step 18) so that they can be later retrieved and correlated with the
20 derived offset watermark vectors (see, e.g., step 38). This step can be largely omitted by the following process.

The original offset watermark vectors are computed using a secret random hash function H . The function H maps copyright and other information that the user desires to embed in the document (e.g., "This picture is the
25 property of XYZ Corp., unauthorized copying is forbidden") to the sequence of numbers $W = w_1, \dots, w_n$ that was used as the original offset watermark vector. The sequence of numbers preferably has same structure and function as discussed above and appear to be random, but the sequence is easily reconstructed given the secret function H and the underlying information to be
30 inserted into the document. Hence, a watermark is identified by reconstructing

the original offset watermark vector locally instead of retrieving the vector from a database. More generally, the text to be embedded may be a simple serial number, and this serial number can be retrieved from the document by checking all possibilities to see if there is a correlation. This check can be
5 done locally if H is available, since all relevant original offset vectors can be regenerated as needed.

Thus, according to this variation of the present invention, one need not subtract the original picture before carrying out the dot product form of the correlation test described above in the main embodiment. In such case, the
10 correlation test generates the old dot product (which is large, precisely what is desired) plus the dot product of the offset vector and the original picture. Since the offset vector is random, this dot product is small (in the noise range) for any picture. Therefore, one does not need the original picture to do the correlation test. Moreover, by using the secret random hash function H , one need not
15 store the offset vectors. The function maps a copyright notice or text into a sequence of independent Gaussian offsets (i.e., an offset vector). Then, one may choose the offset vector for some text to be $H(\text{text})$. Now, one need only remember the text, not the whole offset vector. The text may be timestamped so that the same offset vector is only used once, although one can use the
20 same offset vector more than once.

This method is provably secure, even against colluders, but has low memory requirements. A two-tiered version, wherein there are two hash functions (e.g., one for the sign and one for the magnitude of the offsets) might be used as well. In this way, one of the two (sign or magnitude) would be kept
25 in reserve and not released, even in the secure software. More generally, a series of different watermarks might be used and released according to different purposes. For example, a "Do not copy" watermark might be used where the author is not seeking to restrict "access" whereas a "Do not access" watermark might be used where the author desires to receive payment before

access is allowed. Alternatively, a series of watermarks may be used to restrict the number of accesses.

Another two-tiered approach involves one level of watermark that may be somewhat easy to detect without knowing any secret key and another level
5 that is much more secure but requires a secret key or other secret information to detect.

The above-described variants can be combined advantageously to provide a scheme to prevent unauthorized copying of certain media such as CD's and VCR videos. In this application, a given text -- such as "Do not copy"
10 -- is used as the watermark. A VCR can then check for the presence of this watermark before allowing the copying to take place. This would be achieved by having the secret function H embedded in the VCR software or hardware in a secure fashion, e.g., through a secure chip or via a protected software encryption scheme. The value of H would also be embedded securely in the
15 hardware or software that generates the watermarked copy in the first place instance.

In the VCR/CD application, it may only be necessary to use a single watermark for many copies of the document, in which case it may only be necessary to use a single watermark offset vector (e.g., H ("Do not copy")) for
20 different documents. In this variant, the system must be secure against a different kind of collusion; namely, one in which the same watermark is used with different documents instead of the case where the same document is used with different watermarks. Fortunately, the same analysis applies to both scenarios equally well, such that either scheme is secure against collusion.

25 In the above-described variant, the hardware/software that creates the watermarks is in secure hands (so that H remains secret and cannot be misused). For example, if the adversary is allowed to watermark a blank document, then the scheme can lose security. There are several ways, however, that security can be enhanced as is now explained.

In one approach, it is assumed that each copy of the watermarking software produces watermarks unique to the copy. For example, the XYZ Corporation watermarking software produces watermarks of the form $H(\text{XYZCORP} \mid \text{Do not copy})$. Then, only the watermarks produced by that software would be compromised if the XYZ software were stolen. (For simplicity, each version of the software could be the same except for a special key unique to the version.) Alternatively, the original offset watermark vectors can be derived as a function of the document that is being watermarked in addition to the text that is being embedded into the document. This has the effect of making watermarks corresponding to "Do not copy" be different for each document in which they appear. For example, one might use $H(x_1 \dots x_n \mid \text{Do not copy})$ as the original offset watermark vector for a document with features x_1, \dots, x_n into which the "Do not copy" text is embedded. Even further, the string x_1, \dots, x_n may include random numbers so that offset vectors can be further differentiated in an effort to prevent attacks.

In order to confirm the presence of a watermark in the preceding examples, one still needs to know (or guess, perhaps by exhaustive search) the underlying text that was used to generate the original offset vector. This process can be simplified by embedding serial numbers instead of text. Once the serial number is retrieved, a global database is consulted to find out what the text is. However, it is still necessary to be careful how a serial number is embedded since exhaustive search over a space of 12-digit numbers would be costly and difficult. In such a case, it would be much better to separately embed say four (4) serial numbers, each with 3 digits. (Of course, such numbers and their characteristics are merely exemplary). Then, one would only have to search over a space of 1000 numbers (instead of 1,000,000,000,000 numbers) four times. (This technique makes use of the fact that the watermarking procedures can be used to embed more than one watermark in a document.) One watermark could be used for each decimal or letter in a serial number. As a specific example, if a given letter of a serial

number is "a", and this letter appears in the third position of the number, then the watermark could be a random string computed by generating a hash $H(3,a)$. Alternatively, the $H(a)$ could be used to generate the watermark, which would then be placed in the third component of the picture.

5 It is also possible to make the watermarking process more resilient to noise as well as more secure. This is achieved as follows.

Suppose that one desires to embed the text "Do not copy" in a document. Another good way of doing this is to embed multiple offset watermark vectors in the document. For example, we could use $H(y_1 \mid \text{Do not copy})$, $H(y_2 \mid \text{Do not copy})$, ..., $H(y_m \mid \text{Do not copy})$ for different values of y_1 , ..., y_m as the vectors. If any of the watermarks is detected, then copying would not proceed. Such a scheme is more robust since all m vectors would have to be ruined by noise or be removed by an adversary before copying could proceed. If there is a chance p of being able to remove any one of the
10 vectors, then the chance of losing all m is p^m (assuming independence),
15 which is very small (e.g., if $p=.01$ and $m=4$, then $p^m = 10^{-8}$).

There are several multiple watermarks can be embedded in the document. One method would be to combine the multiple watermarks with the same baseline watermark vector, e.g., by simply adding them all together.
20 Alternatively, each watermark vector could be used with a different baseline vector, e.g., when each watermark is placed in a different component of the document.

Additional variants of the present invention are now described. One variation requires a user to have a password before being able to read or
25 process a document. In particular, when the watermark is generated by a hash function, such as $H(\text{XYZCORP} \mid \text{Do not copy})$, the watermark may be of the form $H(\text{password}^{****} \text{ required for access})$, where $****$ is the password. In this case, processing of the document is allowed only if the watermark is detected (as opposed to the case when processing is not allowed when a watermark is

detected). In this alternate embodiment, the user needs to know the password in order for the watermark (which depends on the password) to be detected.

Yet another variant facilitates tracing of the history of a document. In particular, whenever a person touches or possesses a document, a watermark
5 is added to the document with the ID of that person. In this way, if the document is released illegally, the last person to touch or possess the document can be determined. Moreover, each time a watermark is added, one could also add a timestamp to determine the last possessor's identity.

Another variant is a method to reduce noise in the correlation test
10 (previously described) to thereby decrease the occurrence of false positives and false negatives when checking for a watermark. In this embodiment, some normalization on the baseline watermark and/or the offset watermark is carried out. For example, if the i th component of the baseline watermark x_i is replaced by $x_i + w_i$ in the watermarked document, then the procedure
15 involves several steps that are now described:

(1) The routine computes basic statistics (such as average value and standard deviation) for each x_i . This can be done by generating x_i for an ensemble of documents and taking the mean and standard deviation of the observed values. It could also be done by generating x_i for the single work
20 being protected but from different portions or manifestations of the work. (For example, with a movie, one could compute values for x_i by looking over several frames; with a picture, one could look over several portions of the picture.)

(2) Modify x_i and w_i by normalizing with respect to the statistics. If
25 one computed the mean of x_i , then the routine would subtract this value from the actual x_i . If one computed the standard deviation, then the routine would divide this value into the actual x_i . Alternatively, one could multiply the standard deviation times w_i . More generally, the watermark value can be scaled by an amount derived from analysis of other pictures or regions of the
30 picture being watermarked.

The above processing is useful because it helps all values in the correlation test have substantially equal magnitude and therefore noise (or a particular error) cannot get too much weight.

(3) A further variation is to draw the normalization values from the document itself. This is especially useful in audio or movie applications where one can nearly deduce the original baseline watermark values x_i from the document without having to look them up in a database (because there are often many frames in a movie that are nearly identical). Once one has an approximation to x_i , this approximation can be subtracted before the correlation test is performed (as previously described). In other words, the document contains redundant information that can be used to regenerate the baseline watermark so that it can be subtracted (but without having to look it up in a database).

Thus, one can watermark each frame of a movie separately. To check for a watermark in one frame, one can use the previous frame in place of the original when checking for the watermark, e.g. subtract the previous frame (instead of the original, which is not necessarily available) before doing the correlation test. More generally, this technique can be used whenever one has available a copy of the image that is similar to the original but not identical to the watermarked copy being evaluated.

The present invention also contemplates further variants and/or modifications, which are now also described. One of the features of the invention described above involves adding the watermark vector to the baseline watermark. Although this is desirable, instead of merely adding the watermark vector to the baseline watermark, a more complicated combination may be done. In particular, instead of $x_i + w_i$, one might compute $x_i(1 + w_i)$. More generally, one can replace x_i by a function $f_i(X, W)$. It is also practicable to use scaling, e.g., multiplying the value of w_i by a scalar so that the intensity of the watermark can be adjusted.

Further, the approach of using a watermarked copy of the original image in place of the original image during a correlation test (for a different watermark) can also be used to protect the original image after the test is run. For example, consider the following scenario. An author/creator of a photograph finds a copy of his or her work posted on the World Wide Web, the Internet's multimedia information retrieval system. Assume that the author now desires to prove that it carries one of the author's watermarks. One could reveal the original picture to a judge (or whomever is checking the claim), who would then subtract it from the watermarked copy and run a correlation test with the alleged watermark. The problem with this approach is that the judge has a copy of the original picture without the watermark. If this copy is stolen, the evaluation cannot be run since the copy does not carry a watermark. However, one could improve the process by giving the judge a copy of the image with a different watermark. This image is very close to the original; thus, it will be sufficient for the judge's purposes. Release of the original picture (i.e., without the author's consent), however, will be avoided.

The approach in the previous example can be carried further in ways that should be especially valuable for the "Do not copy" application. As previously described, a secure memory may be used to store or compute the watermark corresponding to "Do not copy." Indeed, all known schemes need to have the watermark that is being checked remain secret. This is because once you have the watermark, it is easy to remove it (e.g., subtract it from the watermarked copy to obtain a clean copy). Unfortunately, the correlation test needs to know the watermark in order to run the correlation test. That is why secure hardware in VCR (for example) is used. However, one can overcome the need for the secure hardware as follows.

Suppose one wants to test for a watermark vector w in a document X' . In the original procedure, one would process X' and then run a correlation test (e.g., by computing the dot product $w * X'$). If w was present in X' , then the correlation would be high. In the alternative embodiment now described, one

does not release w at all. Rather, $(w+w')/2$ is released and the correlation is run against $(w+w')/2$, where w' is another watermark vector. Because watermark vectors can be made to appear random, it is not possible for an adversary to learn anything about w from seeing $(w+w')/2$. However, the correlation test will be positive if and only if w was present. The result of the test will be weaker by a factor of 2, but this is well within tolerance.

In summary, it is possible to run the correlation test without revealing information about the watermark vector, because the watermark vector is "masked" before it is released. The masked watermark vector will still perform well in the correlation test. This idea can be extended by providing each VCR with a different masking of the "Do not copy" watermark vector. Thus if one VCR is compromised, it will not help the adversary remove the watermark for any other VCR. In fact, if the adversary uses his knowledge of $(w+w')/2$ to modify his picture so that the correlation test with $(w+w')/2$ is negative (which he can only do by subtracting a multiple of this vector from the image), he will have unwittingly embedded the new watermark w' in the picture (and he will not have removed the original watermark w). Thus, not only will the original watermark still be present, but there will be proof that the adversary tried to cheat; further, the party will know which VCR was opened for this purpose.

The "Do not copy" text described above is not meant to be taken by way of limitation. Depending on the particular application, other warning(s) can be used, such as "Do not allow access", "Do not allow access unless a password is given", "Do not allow processing", or variations and/or combinations of the above.

As discussed above, it has been suggested that the watermark be placed in a critical region of the data, e.g., in a spatio-temporal frequency domain of the work. One particular advantageous method for achieving this would be use of a spectral transform (e.g., the discrete cosine transform (DCT) or other transforms) to form the baseline watermark of the data. A "critical

region" is a region that, if destroyed, would result in serious degradation of the data.

CLAIMS

1. A method of copy protection for a document, comprising the steps of:

generating a first digital string from the document to form a baseline
5 watermark;
generating a second digital string from given text;
generating a watermark having a predetermined relationship to the first
and second digital strings; and
inserting the watermark into the document to protect the document
10 against illicit copying.

2. The method of copy protection as described in Claim 1 further including the steps of:

retrieving a derived watermark from the document to form a third digital
15 string;
generating a fourth digital string from the given text;
running a correlation test between the third and fourth digital strings;
and
if the third and fourth digital strings have a predetermined correlation,
20 generating an indication that the given text is present in the document.

3. An access control method, comprising the steps of:
during a watermarking phase:

generating a first digital string from an object sought to be
25 protected to form a baseline watermark;
generating a second digital string from given text;
generating a watermark having a predetermined relationship
to the first and second digital strings; and
inserting the watermark into the object to protect access to
30 the object;

upon an access request:

retrieving a derived watermark from the object to form a third digital string;

generating a fourth digital string from the given text;

5 correlating the third and fourth digital strings; and

if the third and fourth digital strings have a predetermined correlation, authorizing access to the object.

4. A method of watermarking an object, comprising the steps of:

10 generating a first digital string from the object to form a baseline watermark;

generating a second digital string from given text providing an indication that a first action with respect to the object is allowed and a second action with respect to the object is prohibited;

15 generating a watermark having a predetermined relationship to the given text and the second digital string; and

inserting the watermark into the object.

5. A method for determining whether an object has a given
20 watermark, comprising the steps of:

processing the object to generate a data string;

correlating the data string with a value that is a function of the given watermark and a second watermark such that information useful in determining the given watermark cannot be obtained from the value; and

if the data string and the value correlate to a predetermined extent, indicating that the object has been watermarked with the given watermark.

6. A method for determining whether a document has a given
5 watermark vector embedded therein, comprising the steps of:

processing the document to generate a data string;

correlating the data string with a value that is a function of the
given watermark vector and a second watermark, wherein the second
watermark masks information about the given watermark;

10 accepting the document as including the given watermark vector if
the data string and the value correlate to a predetermined extent.

7. A access protection method operative in a device having
means for outputting given content, comprising the steps of:

15 retrieving a derived watermark and a derived signal from the given
content;

generating a digital string from the derived signal using a secure
hash function;

correlating the derived watermark and the digital string; and

20 based on a result of the correlating step, taking a given action.

8. A method for authorizing access to given content that has a
given watermark embedded therein, comprising the steps of:

processing the given content to generate a first data string;

generating a second data string by applying a given function to the first data string;

correlating the first and second data strings; and

if the first and second data strings correlate to a given degree,
5 providing a password to enable further processing of the given content.

9. A method for computing a derived watermark, comprising the steps of:

processing a given work W to form a first data string $x_1, x_2, \dots x_n$,

10 processing a collection C of works to form a second data string $y_1, y_2, \dots y_n$; and

computing the derived watermark $z_1, z_2, \dots z_n$ by applying a given function $f(x_i, y_i)$.

15 10. A method of access control for a document, comprising the steps of:

generating a first digital string from the document to form a baseline watermark;

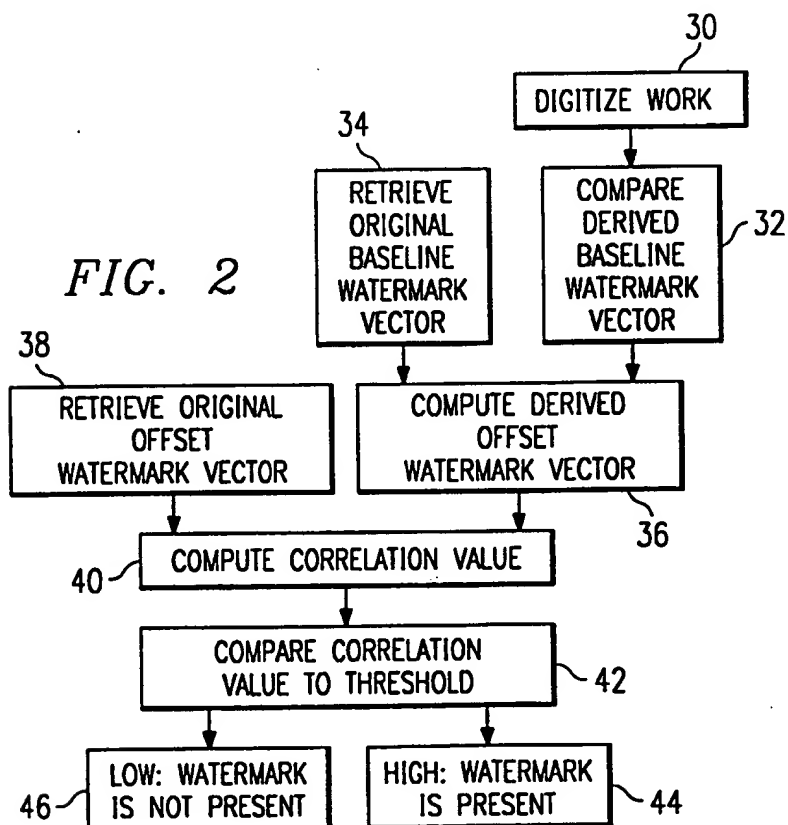
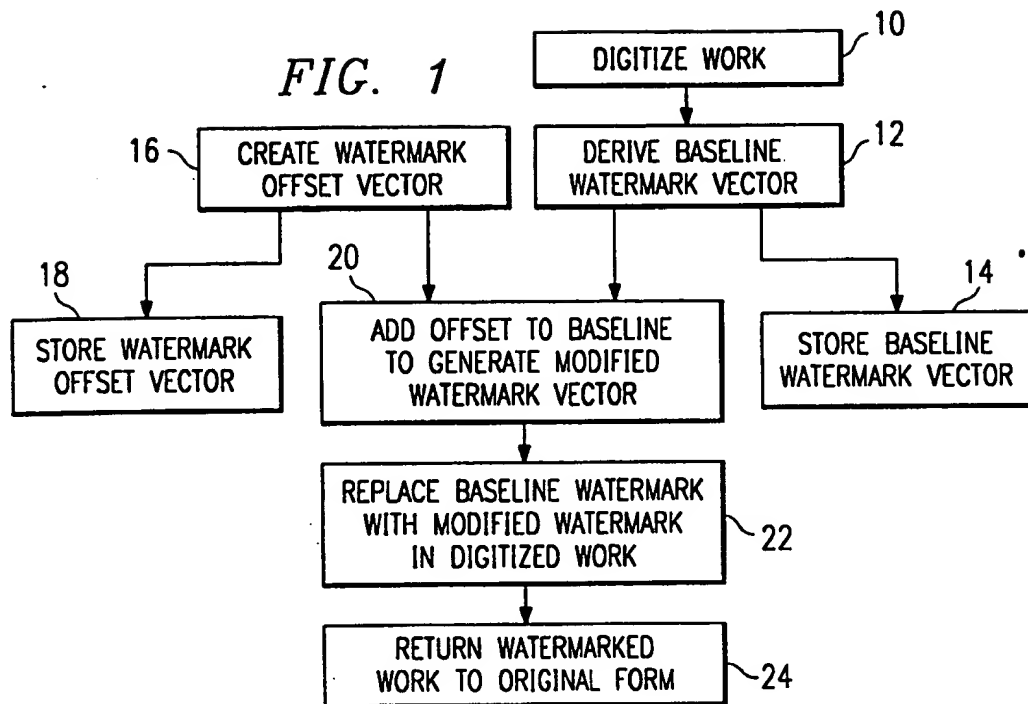
generating a second digital string from given text;

20 generating a set of watermarks each having a predetermined relationship to the first and second digital strings; and

inserting the set of watermarks into the document to protect the document against illicit use.

11. A method for detecting a watermark in a document using information that, if disclosed, does not compromise security of the document, comprising the steps of:

- processing the document to generate a data string;
- 5 correlating the data string with the information; and
- accepting the document as including the watermark if the data string and the information correlate to a predetermined extent.



(51) International Patent Classification ⁶ : H04L 9/00		A3	(11) International Publication Number: WO 99/10858
			(43) International Publication Date: 4 March 1999 (04.03.99)
(21) International Application Number: PCT/US98/17833		(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 28 August 1998 (28.08.98)			
(30) Priority Data: 08/920,365 29 August 1997 (29.08.97) US		Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	
(71)(72) Applicant and Inventor: LEIGHTON, F., Thomson [US/US]; 15 Charlesden Park, Newtonville, MA 02160 (US).		(88) Date of publication of the international search report: 24 June 1999 (24.06.99)	
(74) Agent: JUDSON, David, H.; Hughes & Luce, L.L.P., Suite 2800, 1717 Main Street, Dallas, TX 75201 (US).			

```

graph TD
    10[10 DIGITIZE WORK] --> 12[12 DERIVE BASELINE WATERMARK VECTOR]
    10 --> 16[16 CREATE WATERMARK OFFSET VECTOR]
    12 --> 14[14 STORE BASELINE WATERMARK VECTOR]
    12 --> 20[20 ADD OFFSET TO BASELINE TO GENERATE MODIFIED WATERMARK VECTOR]
    16 --> 18[18 STORE WATERMARK OFFSET VECTOR]
    20 --> 22[22 REPLACE BASELINE WATERMARK WITH MODIFIED WATERMARK IN DIGITIZED WORK]
    22 --> 24[24 RETURN WATERMARKED WORK TO ORIGINAL FORM]
  
```

The work to be protected is digitized (10) and a baseline watermark is derived (12). A watermark offset vector is created (16) and then stored (18). The offset vector is added to the baseline watermark vector to generate a modified watermark vector (20). The baseline watermark vector is replaced with the modified watermark in the digitized work (22). And, finally, the watermarked work is returned to the original form (24).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/17833

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : H04L 9/00

US CL : 380/54

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 380/54, 3, 4, 5, 23, 25, 49, 50, 51, 55, 59; 283/113

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,488,664 A (SHAMIR) 30 January 1996, see abstract.	1-11
A	US 5,530,759 A (BRAUDAWAY et al) 25 June 1996, see abstract.	1-11

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*A* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

19 JANUARY 1999

Date of mailing of the international search report

03 MAY 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

BERNARR EARL GREGORY

Telephone No. (703) 306-4153

